

Penalized Ensemble Kalman Filters for High Dimensional Non-linear Systems

Elizabeth Hou* And Alfred O. Hero

University of Michigan, Ann Arbor, Michigan

Earl Lawrence

Los Alamos National Laboratory, Los Alamos, New Mexico

**Corresponding author address:* University of Michigan, Ann Arbor, MI.
E-mail: emhou@umich.edu

ABSTRACT

The ensemble Kalman filter (EnKF) is a data assimilation technique that uses an ensemble of models, updated with data, to track the time evolution of a non-linear system. It does so by using an empirical approximation to the well-known Kalman filter. Unfortunately, its performance suffers when the ensemble size is smaller than the state space, as is often the case for computationally burdensome models. This scenario means that the empirical estimate of the state covariance is not full rank and possibly quite noisy. To solve this problem in this high dimensional regime, a computationally fast and easy to implement algorithm called the penalized ensemble Kalman filter (PEnKF) is proposed. Under certain conditions, it can be proved that the PEnKF does not require more ensemble members than state dimensions in order to have good performance. Further, the proposed approach does not require special knowledge of the system such as is used by localization methods. These theoretical results are supported with superior performance in simulations of several non-linear and high dimensional systems.

1. Introduction

The Kalman filter is a classic technique to track a linear system over time, and many variants based on the extended and ensemble Kalman filters have been proposed to deal with non-linear systems. The ensemble Kalman filter (EnKF) is particularly popular when the non-linear system is extremely complicated and its gradient is infeasible to calculate, which is often the case in geophysical systems. However, these systems are often high dimensional and forecasting each ensemble member forward through the system is computationally expensive. Thus, the filtering often operates in the high dimensional regime where the number of ensemble members, n , is much less than the number of states, p . It is well known that even when $p/n \rightarrow \text{const.}$ and the samples are from a Gaussian distribution, the eigenvalues and the eigenvectors of the sample covariance matrix do not converge to their population equivalents, Johnstone (2001); Johnstone and Yu Lu (2004). Since our ensemble is both non-Gaussian and high dimensional ($n \ll p$), the sample covariance matrix of the forecast ensemble will be extremely noisy. In this paper, we propose a variant of the EnKF specifically designed to handle covariance estimation in this difficult regime, but with weaker assumptions and less prior information than competing approaches.

Other Work

To deal with the sampling errors, many schemes have been developed to de-noise the forecast sample covariance matrix. These schemes “tune” the matrix by correcting for the eigenvalues with variance inflation and the eigenvectors with localization, Hamill et al. (2001); Houtekamer and Mitchell (2001); Ott et al. (2004); Houtekamer et al. (2005); Wang et al. (2007); Anderson (2007, 2009); Li et al. (2009); Bishop and Hodyss (2009a,b); Houtekamer et al. (2009); Campbell et al. (2010); Greybush et al. (2011); Miyoshi (2011). However, these schemes are often not trivial to implement because they require carefully choosing the inflation factor and using expert knowledge of the true system to set up the localization. Additionally, the EnKF with perturbed observations introduces additional sampling errors due to the perturbation noise’s lack of orthogonality with the ensemble. This has led to the development of matrix factorization versions of the EnKF such as the square root and transform filters, Bishop et al. (2001); Whitaker and Hamill (2002); Tippett et al. (2003); Evensen (2004); Hunt et al. (2007); Nerger et al. (2012); Tödter and Ahrens (2015), which do not perturb the observations and are designed to avoid these additional sampling errors.

The ensemble Kalman filter can be seen as a special case of the particle filter, which uses a Gaussian approximation of the conditional state distribution in order to get a closed form update for the analysis ensemble. While the particle filter does not use this approximation, it also requires an exponential number of particles to avoid filter collapse, Snyder et al. (2008). Recently, there has been significant effort to apply the particle filter to larger scale systems using equal weights, van Leeuwen (2010); Ades and van Leeuwen (2013), and merging it with the ensemble Kalman filter to form hybrid filters, Papadakis et al. (2010); Lei and Bickel (2011); Frei and Künsch (2013a); Nakano (2014); Robert and Künsch (2016).

Most similar to our proposed work, Ueno and Tsuchiya (2009) also propose a method using regularized inverse covariance matrices and they justify the appropriateness of using the inverse space with real weather data. However, their method requires the stronger assumptions of Gaussianity and neighborhood knowledge, and is potentially significantly slower.

Proposed Method

We propose a penalized ensemble Kalman filter (PENKF), which uses an estimator of the forecast covariance whose inverse is sparsity regularized. As the localization approaches effectively dampen or zeros out entries in the covariance, our approach zeros out entries in the inverse covariance. As we shall see, this makes a weaker assumption about the relationship between state variables. Further, our approach does not require anything like localization’s detailed knowledge of which covariances to fix at zero or how much to dampen. Instead, it merely favors sparsity in the inverse covariance. Estimation is at most of the same computational order as the other matrix

operations in the EnKF, so the PEnKF is not significantly slower than the EnKF. Additionally, our method is very easy to implement. In the case when we have measurement noise, but not measurement distortion, we can explicitly show the improvement through theoretical guarantees.

OUTLINE

In Section 2, we explain the assumptions in our high-dimensional system and we give background on the EnKF and ℓ_1 penalized inverse covariance matrices. In Section 3, we give details on how to modify the EnKF to our proposed PEnKF and provide theoretical guarantees on the filter. Section 4 contains the simulation results of the classical Lorenz 96 system and a more complicated system based on modified shallow water equations.

2. Background

In this paper, we consider the scenario of a noisy, non-linear system model $f(\cdot)$, which evolves a vector of unobserved states $\mathbf{x}_t \in \mathbb{R}^p$ through time. We observe a noisy vector \mathbf{y}_t , which is a distortion of \mathbf{x}_t by a function $h(\cdot)$. Both the system noise $\boldsymbol{\omega}_t$ and the observation noise $\boldsymbol{\epsilon}_t$ are independent of the states \mathbf{x}_t and each other. We assume both noises are zero mean Gaussian distributed with known diagonal covariance matrices, \mathbf{Q} and \mathbf{R} . Often, it is assumed that the system does not have noise making $\boldsymbol{\omega}_t$ a zero vector, but for generality we allow $\boldsymbol{\omega}_t$ to be a random vector.

$$\begin{aligned}\mathbf{x}_t &= f(\mathbf{x}_{t-1}) + \boldsymbol{\omega}_t && \text{System Model} \\ \mathbf{y}_t &= h(\mathbf{x}_t) + \boldsymbol{\epsilon}_t && \text{Observation Model}\end{aligned}$$

Since our state space is high-dimensional, we assume each state only interacts with a few other states. For any pair of state variables x_i and x_j , we define their interaction as having non-zero conditional correlation, $\text{Cov}(x_i, x_j | x_{-i,j}) \neq 0$ where $x_{-i,j}$ represents all state variables except x_i and x_j . So if we denote the set, \mathcal{E} , as containing all pairs i, j such that $\text{Cov}(x_i, x_j | x_{-i,j}) \neq 0$, its cardinality $s = |\mathcal{E}|$ is sufficiently small such that $s \ll p^2$. Consequently, the number of pairs of state variables that are conditionally uncorrelated is large. Note that this is a different assumption than the one made by most localization approaches which assume that there are a small number of *unconditionally* correlated pairs, $\text{Cov}(x_i, x_j) \neq 0$. Our assumption is much weaker. This conditional uncorrelation condition manifests as zeros in the inverse covariance matrix of \mathbf{x} , as opposed to zeros in the covariance matrix. In other words, the inverse covariance matrix will have s non-zero off-diagonal entries. We can also quantify the sparsity level as d , which is the maximum number of non-zero off-diagonals in any row, so $d^2 \ll p^2$. Finally, because we do not assume that the state variable interactions are the same for different time points, we allow the set \mathcal{E}_t and its size s_t to change over time.

a. Ensemble Kalman Filter

The standard EnKF algorithm Evensen (1994) is shown in Algorithm 1. At time $t = 0$, n samples are drawn from some distribution, which is often chosen as the standard multivariate normal distribution, if the true initial distribution is unknown, to form an initial ensemble $\mathbf{A} \in \mathbb{R}^{p \times n}$. And, at every time point t , the observations \mathbf{y}_t are perturbed n times with Gaussian white noise, $\boldsymbol{\eta}^j \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$, to form a perturbed observation matrix $\mathbf{D}_t \in \mathbb{R}^{p \times n}$, where $\mathbf{d}_t^j = \mathbf{y}_t + \boldsymbol{\eta}^j$.

Algorithm 1 Ensemble Kalman Filter

Input: $\mathbf{A}, \mathbf{H}, \mathbf{Q}, \mathbf{R}$, and \mathbf{D}_t
where \mathbf{H} is the gradient of $h(\cdot)$
for $t \in \{1, \dots, T\}$ **do**
▷ Evolve each ensemble member forward in time
 $\mathbf{a}_0^j = f(\mathbf{a}^j) + \mathbf{w}^j \quad \forall j \in \{1, \dots, n\}$
where $\mathbf{w}^j \sim N(\mathbf{0}, \mathbf{Q})$
▷ Correct the ensemble with the observations
 $\mathbf{A} = \mathbf{A}_0 + \hat{\mathbf{K}}(\mathbf{D}_t - \mathbf{H}\mathbf{A}_0)$
where $\hat{\mathbf{K}} = \hat{\mathbf{P}}^f \mathbf{H}^T (\mathbf{H} \hat{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})^{-1}$
▷ Predict using the analysis ensemble mean
 $\hat{\mathbf{x}}_t = \frac{1}{n} \sum_{j=1}^n \mathbf{a}^j$
end for
Output: $\hat{\mathbf{x}}_t$

The forecast covariance estimator $\hat{\mathbf{P}}^f$ is typically the sample covariance of the forecast ensemble, defined as $\widehat{\text{Cov}}(\mathbf{A}_0) = \frac{1}{n-1} (\mathbf{A}_0 - \bar{\mathbf{A}}_0)(\mathbf{A}_0 - \bar{\mathbf{A}}_0)^T$, where $\bar{\mathbf{A}}_0$ is the sample mean vector, but it can be another estimator such as a localized estimator, or a penalized estimator as proposed in this paper.

b. Bregman Divergence and the ℓ_1 Penalty

Below, we give a brief overview of the ℓ_1 penalized log-determinant Bregman divergence and some properties of its minimizing estimator, as described in Ravikumar et al. (2011). We denote \mathbf{S} to be any arbitrary sample covariance matrix, and $\Sigma = E(\mathbf{S})$ to be its true covariance matrix, where $E(\cdot)$ is the expectation function.

The Bregman divergence is a very general method to measure the difference between two functions. Here the functions to be compared are covariance matrices. Since we are interested in finding a sparse positive definite estimator for the inverse covariance matrix, a natural choice of Bregman function is $-\log \det(\cdot)$, which has a domain restricted to positive definite matrices. Thus Θ , our optimal estimator for the inverse covariance matrix Σ^{-1} , will minimize

$$\arg \min_{\Theta \in \mathbb{S}_{++}^{p \times p}} -\log \det(\Theta) - \log \det(\Sigma) + \text{tr}(\Sigma(\Theta - \Sigma^{-1}))$$

where $\mathbb{S}_{++}^{p \times p}$ is the set of all symmetric positive definite $p \times p$ matrices. This loss function requires the covariance matrix Σ to be known, but it can be approximated by an empirical loss, which replaces Σ with its empirical equivalent \mathbf{S} and adds a penalty term to ensure strict convexity.

The empirical Bregman divergence with function $-\log \det(\cdot)$ and an ℓ_1 penalty term essentially reduces (by dropping the constants) to

$$\arg \min_{\Theta \in \mathbb{S}_{++}^{p \times p}} -\log \det(\Theta) + \text{tr}(\Theta \mathbf{S}) + \lambda \|\Theta\|_1 \quad (1)$$

where $\lambda \geq 0$ is a penalty parameter, and $\|\cdot\|_1$ denotes an element-wise ℓ_1 norm. Thus we will denote the objective in (1) as minimizing $B^\lambda(\cdot \|\mathbf{S}^{-1})$.

This objective has a unique solution, $\Theta = (\hat{\mathbf{S}})^{-1}$, which satisfies

$$\frac{\partial}{\partial \Theta} B^\lambda(\Theta \|\mathbf{S}^{-1}) = \mathbf{S} - \Theta^{-1} + \lambda \partial \|\Theta\|_1 = 0$$

where $\partial||\Theta||_1$ is a subdifferential of the ℓ_1 norm defined in (A2) in the appendix. The solution $(\tilde{\mathbf{S}})^{-1}$ is a sparse positive definite estimator of the inverse covariance matrix Σ^{-1} , and we can write its inverse explicitly as $\tilde{\mathbf{S}} = \mathbf{S} + \lambda \tilde{\mathbf{Z}}$, where $\tilde{\mathbf{Z}}$ is the unique subdifferential matrix that makes the gradient zero.

Ravikumar et al. (2011) show that for well-conditioned covariances and certain minimum sample sizes, the estimator $(\tilde{\mathbf{S}})^{-1}$ has many nice properties including having, with high probability, the correct zero and signed non-zero entries and a sum of squared error that converges to 0 as $n, p, s \rightarrow \infty$. These properties will allow our method, described in the next section, to attain superior performance over the EnKF.

3. ℓ_1 Penalized Ensemble Kalman Filter

Our penalized ensemble Kalman filter modifies the EnKF, by using a penalized forecast covariance estimator $\tilde{\mathbf{P}}^f$. This penalized estimator is derived from its inverse, which is the minimizer of $B^\lambda(\cdot || (\hat{\mathbf{P}}^f)^{-1})$. Thus from Section 2.b, it can be explicitly written as $\tilde{\mathbf{P}}^f = \hat{\mathbf{P}}^f + \lambda \tilde{\mathbf{Z}}$, implying that we essentially learn a matrix $\tilde{\mathbf{Z}}$, and use it to modify our sample covariance $\hat{\mathbf{P}}^f$. From this, our modified Kalman gain matrix is

$$\tilde{\mathbf{K}} = (\hat{\mathbf{P}}^f + \lambda \tilde{\mathbf{Z}}) \mathbf{H}^T \left(\mathbf{H}(\hat{\mathbf{P}}^f + \lambda \tilde{\mathbf{Z}}) \mathbf{H}^T + \mathbf{R} \right)^{-1}.$$

The intuition behind this estimator is that since only a small number of the state variables in the state vector \mathbf{x}_t are conditionally correlated with each other, the forecast inverse covariance matrix $(\mathbf{P}^f)^{-1}$ will be sparse with many zeros in the off-diagonal entries. Furthermore, since minimizing $B^\lambda(\cdot || (\hat{\mathbf{P}}^f)^{-1})$ gives a sparse estimator for $(\mathbf{P}^f)^{-1}$, this sparse estimator will accurately capture the conditional correlations and uncorrelations of the state variables. Thus $\tilde{\mathbf{P}}^f$ will be a much better estimator of the true forecast covariance matrix \mathbf{P}^f because the ℓ_1 penalty will depress spurious noise in order to make $(\tilde{\mathbf{P}}^f)^{-1}$ sparse, while the inverse of the sample forecast covariance $(\hat{\mathbf{P}}^f)^{-1}$, when it exists, will be non-sparse. As in most penalized estimators, the $\tilde{\mathbf{P}}^f$ is a biased estimator of the forecast covariance, while the sample forecast covariance is not. But because the forecast distribution is corrected for in the analysis step, it is acceptable to take this bias as a trade-off for less variance (sampling errors). Additionally, this bias due to penalization in the inverse covariance behaves in a similar way as variance inflation, so having a biased estimator is not necessarily disadvantageous. Furthermore, since we do not assume the state variables interact in the same way over all time, we re-learn the matrix $\tilde{\mathbf{Z}}$ every time the ensemble is evolved forward.

We can choose the penalty parameter λ in a systematic fashion by calculating a regularization path, solving (1) for a list of decreasing λ s, and evaluating each solution with an information criterion such as an extended or generalized AIC or BIC, Foygel and Drton (2010); Lv and Liu (2014). Additionally, if we have knowledge or make assumptions about the moments of the ensemble's distribution, we know the optimal proportionality of the penalty parameter (see proof of Theorem 3.1). Thus, we can refine the penalty parameter by calculating a regularization path for the constant of the optimal order. In Section 4, we describe a practical approach to choosing λ using a free forecast model run like in Robert and Künsch (2016) and the BIC.

a. Implications on the Kalman Gain Matrix

When $h(\cdot) = \mathbf{H}$ is a diagonal matrix, then the observations do not contain distortions of the interactions between state variables in this case. So the true Kalman gain matrix \mathbf{K} inherits many

of the properties of the forecast covariance matrix \mathbf{P}^f , and our modified Kalman gain matrix $\tilde{\mathbf{K}}$ will benefit from many of the nice properties of our forecast covariance estimator $\tilde{\mathbf{P}}^f$.

How good of an estimator we can get for the forecast covariance matrix \mathbf{P}^f will of course depend on its structure. If it is close to singular or contains lots of entries with magnitudes smaller than the noise level, it will be always be difficult to estimate. So for the following theorem, we assume that the forecast covariance matrix is well-behaved, which we formally state below.

- (i) There exists some $\alpha \in (0, 1]$ such that $\max_{e \in \mathcal{E}^c} \|\Gamma_{e\mathcal{E}}(\Gamma_{\mathcal{E}\mathcal{E}})^{-1}\|_1 \leq (1 - \alpha)$ where Γ is the Hessian of $B^\lambda(\Theta \| (\mathbf{H}\hat{\mathbf{P}}^f\mathbf{H}^T + \mathbf{R})^{-1})$.
- (ii) The ratio between the maximum and minimum eigenvalues of \mathbf{P}^f is bounded.
- (iii) The maximum ℓ_1 norms of the rows of \mathbf{P}^f and $(\Gamma_{\mathcal{E}\mathcal{E}})^{-1}$ are bounded.
- (iv) The minimum non-zero value of $(\mathbf{P}^f)^{-1}$ is $\Omega(\sqrt{\log(p)/n})$ for a sub-Gaussian state vector and $\Omega(\sqrt{p^{3/m}/n})$ for state vectors with bounded $4m^{\text{th}}$ moments.

Theorem 3.1. *Under the above assumptions and for the system described in Section 2 (\mathbf{H} and \mathbf{R} are diagonal matrices), the modified Kalman gain matrix $\tilde{\mathbf{K}}$ has*

- (a) *the same zero and signed non-zero off-diagonal entries as \mathbf{K} , the Kalman gain matrix derived using the true forecast covariance \mathbf{P}^f*
- (b) *a sum of squared error that converges to 0 as the sample size n , dimensionality p , and total number of non-zero off-diagonals entries s , all increase to ∞*

with high probability as long as the sample size is at least $o(n) = d^2 \log(p)$ for sub-Gaussian ensembles and $o(n) = d^2 p^{3/m}$ for ensembles with bounded $4m^{\text{th}}$ moments.

The above theorem tells us the minimum sample size required for our modified Kalman gain matrix to be a good estimate of the true Kalman gain matrix. The sub-Gaussian criterion, where all moments are bounded, is actually very broad and includes any state vectors with a strictly log-concave density and any finite mixture of sub-Gaussian distributions. However even if we cannot bound all the moments, the more $4m^{\text{th}}$ moments we can bound, the smaller a sample size is needed. Notice that as long as we can bound at least 16 moments, we do not need $n > p$. It is also clear that our modified Kalman gain matrix requires many fewer samples to be a good estimator, because the sample Kalman gain matrix requires $o(n) = p^2$ samples in the sub-Gaussian case, and also significantly more in the other case, Vershynin (2012).

The theoretical results above do not mean that when \mathbf{H} is not diagonal, the PEnKF will do poorly. It is simply far more complicated to prove theoretical guarantees because they will depend on the form of \mathbf{H} . Intuitively, since the only randomness in the Kalman gain matrix is the forecast covariance, then as long as we provide a better estimator of it than the sample covariance, we will still have good performance independent of \mathbf{H} .

b. Implications on the Analysis Ensemble

It is well known that due to the additional stochastic noise used to perturb the observations, the covariance of the EnKF's analysis ensemble, $\widehat{\text{Cov}}(\mathbf{A})$ is not equivalent to its analysis covariance calculated by the Gaussian update $\hat{\mathbf{P}}^a = (\mathbf{I} - \hat{\mathbf{K}}\mathbf{H})\hat{\mathbf{P}}^f$. This has led to the development of deterministic variants such as the square root and transform filters, which do have $\widehat{\text{Cov}}(\mathbf{A}) = \hat{\mathbf{P}}^a$. However, in a non-linear system, this update is sub-optimal because it uses a Gaussian approximation of

$\Pr(\mathbf{x}_t|\mathbf{y}_{t-1})$, the actual distribution of forecast ensemble \mathbf{A}_0 . Thus let us denote \mathbf{P}^a as the true analysis covariance defined as

$$\int (\mathbf{x}_t)^2 \Pr(\mathbf{x}_t|\mathbf{y}_t) d\mathbf{x}_t - \left(\int \mathbf{x}_t \Pr(\mathbf{x}_t|\mathbf{y}_t) d\mathbf{x}_t \right)^2$$

where $\Pr(\mathbf{x}_t|\mathbf{y}_t) = \Pr(\mathbf{y}_t|\mathbf{x}_t) \Pr(\mathbf{x}_t|\mathbf{y}_{t-1}) / \Pr(\mathbf{y}_t)$ is not Gaussian. Then, $E(\hat{\mathbf{P}}^a) \neq \mathbf{P}^a$ and there will always be this analysis spread error regardless of whether $\widehat{\text{Cov}}(\mathbf{A}) = \hat{\mathbf{P}}^a$ or not.

As also mentioned in Lei and Bickel (2011), actually none of the analysis moments of the EnKF are consistent with the true moments including the analysis mean. However this analysis error is present in all methods that do not introduce particle filter properties to the EnKF, and thus is not the focus of our paper. We are primarily concerned with the sampling errors in high-dimensional systems and simply wanted to address that the lack of equivalence to the Gaussian update is irrelevant in our case of a non-linear system.

c. Computational Time and Storage Issues

The computational complexity of solving for the minimizer of $B^\lambda(\cdot|\|\hat{\mathbf{P}}^f)$ with the GLASSO algorithm from Friedman et al. (2008) is $O(sp^2)$, which is, at most, of the same order as matrix inversion. So our PEnKF algorithm will not take considerably more time than the standard EnKF algorithm. Additionally when \mathbf{H} is a diagonal matrix, we know from Lemma A0.2, that the Kalman gain matrix \mathbf{K} is sparse. Thus we can directly solve for a sparse minimizer Θ of $B^\lambda(\Theta|\|(\mathbf{H} \frac{1}{n-1}(\mathbf{A}_0 - \hat{\mathbf{A}}_0)(\mathbf{A}_0 - \hat{\mathbf{A}}_0)^T \mathbf{H}^T + \mathbf{R})^{-1})$ by cleverly applying a primal algorithm such as P-GLASSO of Mazumder and Hastie (2012). This can make it possible to remove the necessity of storing a $p \times p$ matrix because our modified Kalman gain matrix $\tilde{\mathbf{K}} = \mathbf{H}^{-1} - \mathbf{H}^{-1} \mathbf{R} \Theta$ can be stored in a sparse matrix structure. This application of the P-GLASSO algorithm will be significantly slower than the GLASSO though, so it should only be used when it is impossible to store $p \times p$ matrices.

4. Simulations

In all simulations, we compare to an ensemble Kalman filter where the forecast covariance matrix is localized with a taper matrix generated from equation (4.10) in Gaspari and Cohn (1999). The taper matrix parameter c is chosen using the true interactions of the system, so the localization should be close to optimal for simple systems. We use this TAPER-EnKF as the baseline because if the PEnKF can do as well as this filter, it implies that the PEnKF can learn a close to optimal covariance matrix, even without the need to impose a known neighborhood structure. If the PEnKF can do better than this filter, it implies that the PEnKF is learning some structure that is not captured by localization with a taper matrix.

In order to choose the penalty parameter for PEnKF, we assume that the state variables in our examples are sub-Gaussian. In this case, we can set $\lambda = c_\lambda \sqrt{R \log(p)/n}$ for some appropriate choice of c_λ (see the proof of Theorem 3.1). To estimate c_λ , we generate a representative ensemble (which may also be our initial ensemble) using a free forecast run like in Robert and Künsch (2016) in which a state vector is drawn at random (e.g. from $N(\mathbf{0}, \mathbf{I})$) and evolved forward. The representative ensemble is produced by taking a set of equally spaced points (e.g. every 100th state vector) from the evolution. This ensemble is used to choose c_λ from some predefined interval by minimizing the eBIC of Foygel and Drton (2010) if $p > n$ or the BIC of Schwarz (1978) if $p < n$.

Of course (1) is not a likelihood unless the states are Gaussian. So, we have a misspecified model where we are treating the states as having a Gaussian likelihood when evaluating a potential penalty parameter using an information criteria. In this case, we should correct our information criterion for the misspecification as in Lv and Liu (2014). However, this can be quite difficult and

we leave an in-depth exploration of this problem for future work. In the meantime, this choice seems to perform well despite its lack of optimality.

We define the root mean squared error (RMSE) used to evaluate a filters performance by

$$\text{RMSE}_t = \sqrt{(\|\hat{\mathbf{x}}_t - \mathbf{x}_t\|_2)^2 / p}$$

where RMSE_t is an element of a vector indicating the RMSE at time point t , \mathbf{x}_t is a vector of the true hidden state variables, $\hat{\mathbf{x}}_t$ is a filter's estimators for the true state vector, and $\|\cdot\|_2$ is the ℓ_2 norm. We will refer to quantiles such as the mean or median RMSE to be the mean or median of the elements of the RMSE vector.

a. Lorenz 96 System

The 40-state Lorenz 96 model is one of the most common systems used to evaluate ensemble Kalman filters. The state variables are governed by the following differential equations

$$\frac{dx_t^i}{dt} = (x_t^{i+1} - x_t^{i-2})x_t^{i-1} - x_t^i + 8 \quad \forall i = 1, \dots, 40$$

where $x_t^{41} = x_t^1$, $x_t^0 = x_t^{40}$, and $x_t^{-1} = x_t^{39}$.

We use the following simulation settings. We have observations for the odd state variables, so $\mathbf{y}_t = \mathbf{H}\mathbf{x}_t + \boldsymbol{\epsilon}_t$ where \mathbf{H} is a 20×40 matrix with ones at entries $\{i, j = 2i - 1\}$ and zeros everywhere else and $\boldsymbol{\epsilon}_t$ is a 20×1 vector drawn from a $N(\mathbf{0}, 0.5\mathbf{I})$. We initialize the true state vector from a $N(\mathbf{0}, \mathbf{I})$ and we assimilate at every $0.4t$ time steps, where $t = 1, \dots, 2000$. The system is numerically integrated with a 4th order Runge-Kutta method and a step size of 0.01. The main difficulties of this system are the large assimilation time step of 0.4, which makes it significantly non-linear, and the lack of observations for the even state variables.

Since the exact equations of the Lorenz 96 model are fairly simple, it is clear how the state variables interact with each other. This makes it possible to localize with a taper matrix that is almost optimal by using the Lorenz 96 equations to choose a half-length parameter c . However, we do not incorporate this information in the PEnKF algorithm, which instead learns interactions by essentially extracting it from the sample covariance matrix. We set the penalty parameter $\lambda = c_\lambda \sqrt{0.5 \log(p)/n}$ by using an offline free forecast run to search for the constant c_λ in the range $[0.1, 10]$ as described at the beginning of this section.

We average the PEnKF estimator of the forecast inverse covariance matrix at the time points 500, 1000, 1500, and 2000 for 50 trials with 25 ensemble members, and we compare it to the "true" inverse covariance matrix, which is calculated by moving an ensemble of size 2000 through time. In Figure 1, each line represents the averaged normalized rows of an inverse covariance matrix and the lines are centered at the diagonal. The penalized inverse covariance matrix does a qualitatively good job of capturing the neighborhood information and successfully identifies that any state variables far away from state variable i , do not interact with it.

Because the PEnKF is successful at estimating the structure of the inverse covariance matrix and thus the forecast covariance matrix, we expect it will have good performance for estimating the true state variables. We compare the PEnKF to the TAPER-EnKF and other estimators from Bengtsson et al. (2003); Lei and Bickel (2011); Frei and Künsch (2013a,b) by looking at statistics of the RMSE. Note that in order to have comparable statistics to as many other papers as possible, we do not add variance inflation to the TAPER-EnKF (like in Bengtsson et al. (2003); Frei and Künsch (2013a,b) and unlike in Lei and Bickel (2011)). Also, like in those papers, we initialize the ensemble from a $N(\mathbf{0}, \mathbf{I})$, and we use this ensemble to start the filters. Note that in this case, the initial ensemble is different than the offline ensemble that we use to estimate the PEnKF's penalty parameter. This is because the initial ensemble is not representative of the system and its sample covariance is an estimator for the identity matrix. The TAPER-EnKF, which is simply

called the EnKF in the other papers, is localized by applying a taper matrix where $c = 10$ to the sample covariance matrix.

We show the mean, median, 10%, and 90% quantiles of the RMSE averaged over 50 independent trials for ensembles of size 400, 100, 25, and 10 in Table 1. For 400 ensemble members, the PEnKF does considerably better than the TAPER-EnKF and its relative improvement is larger than that of the XEnKF reported in Bengtsson et al. (2003) and similar to those of the NLEAF, EnKPF, and XEnKF reported in Lei and Bickel (2011); Frei and Künsch (2013a,b) respectively. For 100 ensemble members, the PEnKF does do worse than the TAPER-EnKF and EnKPF of Frei and Künsch (2013a); this we suspect may be due to the bias-variance trade-off when estimating the forecast covariance matrix. The PEnKF has the most significant improvement over the TAPER-EnKF in the most realistic regime where we have fewer ensemble members than state variables. For both 25 and 10 ensemble members, the PEnKF does considerably better than the TAPER-EnKF and it does not suffer from filter divergence, which Frei and Künsch (2013a) report occurs for the PEnKF at 50 ensemble members.

While it is clear the PEnKF does well even when there are fewer ensemble members than state variables, 40 variables is not enough for the problem to be considered truly high-dimensional. We now consider simulation settings where we increase the dimension of the state space p while holding the number of ensemble members n constant. We initialize the ensemble from the free forecast run and set λ and the taper matrix in the same way as in the previous simulations. We examine the mean RMSE averaged over 50 trials and its approximate 95% confidence intervals in the Figure 2. The mean RMSE of the PEnKF is significantly smaller than the mean RMSE of the TAPER-EnKF for all p . Additionally the confidence intervals of the mean RMSE are much narrower than the ones for the TAPER-EnKF. This suggests that there is little variability in the PEnKF’s performance, while the TAPER-EnKF’s performance is more dependent on the trial, with some trials being “easier” for the TAPER-EnKF than others.

b. Modified Shallow Water Equations System

While the Lorenz 96 system shows that the PEnKF has strong performance because it is successful at reducing the sampling errors and capable of learning the interactions between state variables, the system is not very realistic in that all state variables are identical and the relationship between state variables is very simplistic. We now consider a system based on the modified shallow water equations of Würsch and Craig (2014), which models cloud convection with fluid dynamics equations, but is substantially computationally less expensive than actual numerical weather prediction models. The system has three types of state variables: fluid height, rain content, and horizontal wind speed.

To generate this system we use the R package “modifiedSWEQ” created by Robert (2014), and the same simulation settings as in Robert and Künsch (2016). So we always observe the rain content, but wind speed is only observed at locations where it is raining and fluid height is never observed. Explicitly for the R function *generate.xy()*, we use $h_c = 90.02$, $h_r = 90.4$ for the cloud and rainwater thresholds, a 0.005 rain threshold, $\sigma_r = 0.1$, $\sigma_u = 0.0025$ to be the standard deviation of the observation noise for rain and wind respectively, and $\mathbf{R} = \text{diag}([R_r^2 = 0.025^2 \ R_u^2 = \sigma_u^2])$ to be the estimated diagonal noise covariance matrix. All other parameters are just the default ones in the function. The initial ensemble is drawn from a free forecast run with 10000/60 time-steps between each ensemble member. We give a snapshot of the system at a random time point in Figure 3. There are $p = 300$ state variables for each type, making the state space have 900 dimensions and we assimilate the system every 5 seconds for a total time period of 6 hours. Like in Robert and Künsch (2016), we choose to use only 50 ensemble members and we do not perturb rain observations that are 0, because at these points there is no measurement noise.

The TAPER-EnKF uses a $3p \times 3p$ taper matrix with $c = 5$, however the entries off the $p \times p$ block diagonals are depressed (they are multiplied by 0.9). The NAIVE-LEnKPF uses the same settings as in Robert and Künsch (2016), so a localization parameter of 5km, which gives the same taper matrix as the one used in the TAPER-EnKF, and an adaptive γ parameter. For the PEnKF,

we set the penalty parameter to be a $3p \times 3p$ matrix, $\Lambda = c_\lambda \sqrt{\lambda_R \lambda_R^T \log(3p)/n}$, where the first, second, and last p entries of the vector λ_R are to scale for the perturbation noise of the different state types. So they are 1 for fluid height, R_u for wind, and R_r for rain. We choose the constant c_λ with eBIC like before and search in the range $[.005, 1]$.

Figure 4 shows the mean and approximate 95% confidence intervals of the RMSE for fluid height, wind speed, and rain content over 6 hours of time using 50 trials. The mean RMSE for all three filters are well within each others' confidence intervals for the fluid height and wind variables. For the rain variables, the mean RMSE of neither the TAPER-EnKF nor the NAIVE-LEnKPF are in the PEnKF's confidence intervals and the mean RMSE of the PEnKF is on the boundary of the other two models' confidence intervals. This strongly suggests that the PEnKF's rain error is statistically smaller than the rain errors of the other two filters. Since this simulation is not as simple as the previous ones, the interactions between the state variables are most likely not as effectively captured by the taper matrix or other localization methods, and the results from this simulation suggest that the PEnKF is learning more accurate interactions for the rain variables. We do not show the results of the BLOCK-LEnKPF of Robert and Künsch (2016) because the algorithm suffered from filter divergence in 27 of the 50 trials, and in the trials where it did not fail, it performed very similar to the NAIVE-LEnKPF.

5. Discussion

We propose a new algorithm based on the ensemble Kalman filter that is designed for superior performance in non-linear high dimensional systems. This algorithm we call the penalized ensemble Kalman filter because it uses the popular statistical concept of penalization/regularization in order to make the problem of estimating the forecast covariance matrix well-defined (strictly convex). This in turn both decreases the sampling errors in the forecast covariance estimator by trading it off for bias and prevents filter divergence by ensuring that the estimator is positive definite. The PEnKF is computationally efficient in that it is not significantly slower than the standard EnKF algorithms and easy to implement since it only adds one additional step, and this step uses the well-established GLASSO algorithm available in almost any scientific computing language. We give theoretical results that prove that the Kalman gain matrix constructed from this estimator will converge to the population Kalman gain matrix under the non-simplistic asymptotic case of high-dimensional scaling, where the sample size and the dimensionality increase to infinity.

Through simulations, we show that the PEnKF can do at least as well as, and sometimes better than, localized filters that use much more prior information. We emphasize that by doing just as well as the TAPER-EnKF which has a close to optimal taper matrix, the PEnKF is effectively correctly learning the structure of interactions between the state variables. In a non-simulation setting where there is no ground-truth knowledge of the interactions between state variables, correct localization is much more difficult, making any localized filter's performance likely sub-optimal. In contrast, since the PEnKF does not use any of this "oracle" information, its performance will not differ in this way between simulations and real-life situations. The more complicated simulation, based on the modified shallow water equations, highlights this advantage of the PEnKF through its substantial superior performance in estimating the hidden states of the rain variables. Another feature of the approach is that it seems to require less variance inflation. None was applied to any algorithm in our comparison, but the PEnKF approach never collapsed. The penalization of the inverse covariance actually produces a slight inflation on the diagonal of the covariance, which seems to help in this regard.

While we display a very naive way of searching for a good penalty parameter for the PEnKF in the simulations, it is theoretically incorrect and thus not a way to chose the truly optimal penalty parameter. We do believe deriving a specific information criterion for our PEnKF with correct theoretical properties is very important since the PEnKF can be sensitive to the penalty parameter. However, this model selection in misspecified models problem is not trivial to solve and an

active topic in current statistical research. Therefore, we will leave deriving a theoretically correct information criterion for future work.

Acknowledgments. This work was partially supported by the Consortium for Verification Technology under Department of Energy National Nuclear Security Administration award number de-na0002534 and partially supported by the Laboratory Directed Research and Development program at Los Alamos National Laboratory under project 20150033DR, SHIELDS: Space Hazards Induced near Earth by Large Dynamic Storms - Understanding, Modeling, and Predicting.

APPENDIX

For the following proofs we are assuming the system is as described in Section 2; so the true forecast inverse covariance matrix of the system $(\mathbf{P}^f)^{-1}$ is sparse at every time point t and the observation noise covariance matrix \mathbf{R} is diagonal. Additionally we assume the measurement distortion function $h(\cdot)$ is a diagonal matrix \mathbf{H} .

Lemma A0.1. *From the assumptions above, the inverse covariance matrix of the observations $(\text{Cov}(\mathbf{y}))^{-1} = (\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}$, has the same sparsity pattern as the inverse covariance matrix of the unobserved states $(\text{Cov}(\mathbf{x}))^{-1} = (\mathbf{P}^f)^{-1}$.*

Proof of Lemma A0.1. If $\text{Cov}(x_i, x_j | x_{-i,j}) \neq 0$, then $\text{Cov}(y_i, y_j | x_{-i,j}) = \text{Cov}(h_{ii}x_i + \varepsilon_i, h_{jj}x_j + \varepsilon_j | x_{-i,j}) = h_{ii}h_{jj}\text{Cov}(x_i, x_j | x_{-i,j}) \neq 0$ for all i and j . A similar argument holds for when $\text{Cov}(x_i, x_j | x_{-i,j}) = 0$ and all $i \neq j$. \square

Lemma A0.2. *From the assumptions above, the Kalman gain matrix \mathbf{K} is sparse with s non-zero off-diagonal entries.*

Proof of Lemma A0.2.

$$\begin{aligned} \mathbf{K} &= \mathbf{P}^f\mathbf{H}^T(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1} = \mathbf{H}^{-1}\mathbf{H}\mathbf{P}^f\mathbf{H}^T((\mathbf{H}\mathbf{P}^f\mathbf{H}^T)^{-1} \\ &\quad - (\mathbf{H}\mathbf{P}^f\mathbf{H}^T)^{-1}(\mathbf{R}^{-1} + (\mathbf{H}\mathbf{P}^f\mathbf{H}^T)^{-1})^{-1}(\mathbf{H}\mathbf{P}^f\mathbf{H}^T)^{-1}) \\ &= \mathbf{H}^{-1}\left(\mathbf{I} - (\mathbf{H}\mathbf{P}^f\mathbf{H}^T(\mathbf{R}^{-1} + (\mathbf{H}\mathbf{P}^f\mathbf{H}^T)^{-1}))^{-1}\right) \\ &= \mathbf{H}^{-1} - \mathbf{H}^{-1}(\mathbf{H}\mathbf{P}^f\mathbf{H}^T\mathbf{R}^{-1} + \mathbf{I})^{-1} \\ &= \mathbf{H}^{-1} - \mathbf{H}^{-1}\mathbf{R}(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1} \end{aligned}$$

From Lemma A0.1, $(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}$ is sparse with s non-zero off-diagonal entries. So since \mathbf{H}^{-1} and \mathbf{R} are diagonal, the matrix $\mathbf{K} = \mathbf{H}^{-1} - \mathbf{H}^{-1}\mathbf{R}(\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R})^{-1}$ has the same sparsity pattern in the off-diagonals entries. \square

Lemma A0.3. *Under Lemma A0.1, the estimator $(\mathbf{H}\tilde{\mathbf{P}}^f\mathbf{H}^T + \mathbf{R})^{-1}$ minimizes $\mathbf{B}^\lambda(\cdot || (\mathbf{H}\hat{\mathbf{P}}^f\mathbf{H}^T + \mathbf{R})^{-1})$, where $\tilde{\mathbf{P}}^f$ is the minimizer to $\mathbf{B}^\lambda(\cdot || (\hat{\mathbf{P}}^f)^{-1})$.*

Proof of Lemma A0.3.

$$\begin{aligned}
& \arg \min_{\Theta \in \mathbb{S}_{++}^{p \times p}} B^\lambda(\Theta \| (\mathbf{H}\hat{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})^{-1}) \\
&= \arg \min_{\Theta \in \mathbb{S}_{++}^{p \times p}} \text{tr}(\Theta(\mathbf{H}\hat{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})) - \log \det(\Theta) + \lambda \|\Theta\|_1 \\
&\text{Thus its gradient is} \\
&\mathbf{H}\hat{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R} - \Theta^{-1} + \lambda \partial \|\Theta\|_1
\end{aligned} \tag{A1}$$

where $\partial \|\Theta\|_1$ is the sub-differential of the ℓ_1 norm evaluated at some Θ . Since the derivative of an absolute value is undefined at zero, $\partial \|\Theta\|_1$ can be any number between -1 and 1, when Θ is 0. So, it is the set of all matrices $\mathbf{Z} \in \mathbb{S}^{p \times p}$ such that

$$\mathbf{Z}_{ij} = \begin{cases} \text{sign}(\Theta_{ij}) & \text{if } \Theta_{ij} \neq 0 \\ \in [-1, 1] & \text{if } \Theta_{ij} = 0. \end{cases} \tag{A2}$$

Since we know $\hat{\mathbf{P}}^f - \Theta^{-1} + \lambda \partial \|\Theta\|_1 = 0$ when $\Theta = (\tilde{\mathbf{P}}^f)^{-1}$, the gradient in (A1) equals 0 when $\Theta = (\mathbf{H}\tilde{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})^{-1}$. \square

Proof of Theorem 3.1. The following assumptions are necessary for the minimizer of (1) to have good theoretical properties, Ravikumar et al. (2011). Thus we assume they are true for the theorem.

- (i) There exists some $\alpha \in (0, 1]$ such that $\max_{e \in \mathcal{E}} \|\Gamma_{e\mathcal{E}}(\Gamma_{\mathcal{E}\mathcal{E}})^{-1}\|_1 \leq (1 - \alpha)$ where Γ is the Hessian of $B^\lambda(\Theta \| (\mathbf{H}\hat{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})^{-1})$.
- (ii) The ratio between the maximum and minimum eigenvalues of \mathbf{P}^f is bounded.
- (iii) The maximum ℓ_1 norms of the rows of \mathbf{P}^f and $(\Gamma_{\mathcal{E}\mathcal{E}})^{-1}$ are bounded.
- (iv) The minimum non-zero value of $(\mathbf{P}^f)^{-1}$ is $\Omega(\sqrt{\log(p)/n})$ for a sub-Gaussian state vector and $\Omega(\sqrt{p^{3/m}/n})$ for state vectors with bounded $4m^{\text{th}}$ moments.

Our assumptions are stronger than necessary, and it is common to allow the error rates to depend on the bounding constants above, but for simplicity we give the error rates only as a function of the dimensionality n, p and sparsity s, d parameters.

From Lemmas A0.1 and A0.3, and Sections 3.3 and 3.5 of Ravikumar et al. (2011), we know that for sample size i) $n = \Omega(d^2 \log(p))$ for sub-Gaussian random variables and ii) $n = \Omega(d^2 p^{3/m})$ for variables with bounded $4m^{\text{th}}$ moments, the following hold:

- (a) all entries of $(\mathbf{H}\tilde{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})^{-1}$ have the same zeros and signed non-zeros as $(\mathbf{H}\mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1}$
- (b) $\|(\mathbf{H}\tilde{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})^{-1} - (\mathbf{H}\mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1}\|_F^2 = \begin{cases} i) O((s+p) \log(p)/n) & \text{for } \lambda \asymp \sqrt{\log(p)/n} \\ ii) O((s+p) p^{3/m}/n) & \text{for } \lambda \asymp \sqrt{p^{3/m}/n} \end{cases}$

with probability greater than $1 - 1/p$.

Thus from Lemma A0.2, the off-diagonals of $\tilde{\mathbf{K}} = \mathbf{H}^{-1} - \mathbf{H}^{-1} \mathbf{R} (\mathbf{H}\tilde{\mathbf{P}}^f \mathbf{H}^T + \mathbf{R})^{-1}$ have the same zeros and signed non-zeros as $\mathbf{K} = \mathbf{H}^{-1} - \mathbf{H}^{-1} \mathbf{R} (\mathbf{H}\mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1}$. And, since \mathbf{H} and \mathbf{R} are non-random matrices, the sum of squared error $\|\tilde{\mathbf{K}} - \mathbf{K}\|_F^2$ has the same order as (b). \square

References

- Ades, M., and P. J. van Leeuwen, 2013: An exploration of the equivalent weights particle filter. *Quarterly Journal of the Royal Meteorological Society*, **139** (672), 820–840, doi:10.1002/qj.1995.
- Anderson, J. L., 2007: An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus A*, **59** (2), 210–224, doi:10.1111/j.1600-0870.2006.00216.x, URL <http://dx.doi.org/10.1111/j.1600-0870.2006.00216.x>.
- Anderson, J. L., 2009: Spatially and temporally varying adaptive covariance inflation for ensemble filters. *Tellus A*, **61** (1), 72–83, doi:10.1111/j.1600-0870.2008.00361.x, URL <http://dx.doi.org/10.1111/j.1600-0870.2008.00361.x>.
- Bengtsson, T., C. Snyder, and D. Nychka, 2003: Toward a nonlinear ensemble filter for high-dimensional systems. *Journal of Geophysical Research: Atmospheres*, **108** (D24).
- Bishop, C., and D. Hodyss, 2009a: Ensemble covariances adaptively localized with eco-rap. part 1: tests on simple error models. *Tellus A*, **61** (1).
- Bishop, C. H., B. J. Etherton, and S. J. Majumdar, 2001: Adaptive sampling with the ensemble transform kalman filter. part i: Theoretical aspects. *Monthly Weather Review*, **129** (3), 420–436, doi:10.1175/1520-0493(2001)129<0420:ASWTET>2.0.CO;2.
- Bishop, C. H., and D. Hodyss, 2009b: Ensemble covariances adaptively localized with eco-rap. part 2: a strategy for the atmosphere. *Tellus A*, **61** (1), 97–111, doi:10.1111/j.1600-0870.2008.00372.x.
- Campbell, W. F., C. H. Bishop, and D. Hodyss, 2010: Vertical covariance localization for satellite radiances in ensemble kalman filters. *Monthly Weather Review*, **138** (1), 282–290, doi:10.1175/2009MWR3017.1.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans*, **99** (C5), 10 143–10 162.
- Evensen, G., 2004: Sampling strategies and square root analysis schemes for the enkf. *Ocean dynamics*, **54** (6), 539–560.
- Foygel, R., and M. Drton, 2010: Extended bayesian information criteria for gaussian graphical models. *Advances in neural information processing systems*, 604–612.
- Frei, M., and H. R. Künsch, 2013a: Bridging the ensemble kalman and particle filters. *Biometrika*, **100** (4), 781–800.
- Frei, M., and H. R. Künsch, 2013b: Mixture ensemble kalman filters. *Computational Statistics & Data Analysis*, **58**, 127–138.
- Friedman, J., T. Hastie, and R. Tibshirani, 2008: Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, **9** (3), 432–441.
- Gaspari, G., and S. E. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, **125** (554), 723–757.
- Greybush, S. J., E. Kalnay, T. Miyoshi, K. Ide, and B. R. Hunt, 2011: Balance and ensemble kalman filter localization techniques. *Monthly Weather Review*, **139** (2), 511–522, doi:10.1175/2010MWR3328.1.

- Hamill, T. M., J. S. Whitaker, and C. Snyder, 2001: Distance-dependent filtering of background error covariance estimates in an ensemble kalman filter. *Monthly Weather Review*, **129** (11), 2776–2790.
- Houtekamer, P. L., and H. L. Mitchell, 2001: A sequential ensemble kalman filter for atmospheric data assimilation. *Monthly Weather Review*, **129** (1), 123–137.
- Houtekamer, P. L., H. L. Mitchell, and X. Deng, 2009: Model error representation in an operational ensemble kalman filter. *Monthly Weather Review*, **137** (7), 2126–2143, doi:10.1175/2008MWR2737.1.
- Houtekamer, P. L., H. L. Mitchell, G. Pellerin, M. Buehner, M. Charron, L. Spacek, and B. Hansen, 2005: Atmospheric data assimilation with an ensemble kalman filter: Results with real observations. *Monthly Weather Review*, **133** (3), 604–620, doi:10.1175/MWR-2864.1.
- Hunt, B. R., E. J. Kostelich, and I. Szunyogh, 2007: Efficient data assimilation for spatiotemporal chaos: a local ensemble transform kalman filter. *Physica D: Nonlinear Phenomena*, **230** (1-2), 112–126.
- Johnstone, I. M., 2001: On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, **29** (2), 295–327.
- Johnstone, I. M., and A. Yu Lu, 2004: Sparse principle component analysis. *Unpublished Manuscript*.
- Lei, J., and P. Bickel, 2011: A moment matching ensemble filter for nonlinear non-gaussian data assimilation. *Monthly Weather Review*, **139** (12), 3964–3973.
- Li, H., E. Kalnay, and T. Miyoshi, 2009: Simultaneous estimation of covariance inflation and observation errors within an ensemble kalman filter. *Quarterly Journal of the Royal Meteorological Society*, **135** (639), 523–533, doi:10.1002/qj.371.
- Lv, J., and J. S. Liu, 2014: Model selection principles in misspecified models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **76** (1), 141–167.
- Mazumder, R., and T. Hastie, 2012: The graphical lasso: New insights and alternatives. *Electronic journal of statistics*, **6**, 2125.
- Miyoshi, T., 2011: The gaussian approach to adaptive covariance inflation and its implementation with the local ensemble transform kalman filter. *Monthly Weather Review*, **139** (5), 1519–1535, doi:10.1175/2010MWR3570.1.
- Nakano, S., 2014: Hybrid algorithm of ensemble transform and importance sampling for assimilation of non-gaussian observations. *Tellus A*, **66** (0).
- Nerger, L., T. Janji, J. Schrter, and W. Hiller, 2012: A unification of ensemble square root kalman filters. *Monthly Weather Review*, **140** (7), 2335–2345, doi:10.1175/MWR-D-11-00102.1.
- Ott, E., and Coauthors, 2004: A local ensemble kalman filter for atmospheric data assimilation. *Tellus A*, **56** (5), 415–428.
- Papadakis, N., E. M  min, A. Cuzol, and N. Gengembre, 2010: Data assimilation with the weighted ensemble kalman filter. *Tellus A*, **62** (5), 673–697, doi:10.1111/j.1600-0870.2010.00461.x.
- Ravikumar, P., M. J. Wainwright, G. Raskutti, B. Yu, and Coauthors, 2011: High-dimensional covariance estimation by minimizing 1-penalized log-determinant divergence. *Electronic Journal of Statistics*, **5**, 935–980.

- Robert, S., 2014: *modifiedSWEQ: Simplified cumulus convection with the modified SWEQ*. R package version 0.1, <https://github.com/robertsy/modifiedSWEQ>.
- Robert, S., and H. R. Künsch, 2016: Local Ensemble Kalman Particle Filters for efficient data assimilation. *ArXiv e-prints*, 1605.05476.
- Schwarz, G., 1978: Estimating the dimension of a model. *Ann. Statist.*, **6** (2), 461–464, doi:10.1214/aos/1176344136.
- Snyder, C., T. Bengtsson, P. Bickel, and J. Anderson, 2008: Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, **136** (12), 4629–4640, doi:10.1175/2008MWR2529.1.
- Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker, 2003: Ensemble square root filters. *Monthly Weather Review*, **131** (7), 1485–1490, doi:10.1175/1520-0493(2003)131<1485:ESRF>2.0.CO;2.
- Tödter, J., and B. Ahrens, 2015: A second-order exact ensemble square root filter for nonlinear data assimilation. *Monthly Weather Review*, **143** (4), 1347–1367, doi:10.1175/MWR-D-14-00108.1.
- Ueno, G., and T. Tsuchiya, 2009: Covariance regularization in inverse space. *Quarterly Journal of the Royal Meteorological Society*, **135** (642), 1133–1156.
- van Leeuwen, P. J., 2010: Nonlinear data assimilation in geosciences: an extremely efficient particle filter. *Quarterly Journal of the Royal Meteorological Society*, **136** (653), 1991–1999, doi:10.1002/qj.699.
- Vershynin, R., 2012: How close is the sample covariance matrix to the actual covariance matrix? *Journal of Theoretical Probability*, **25** (3), 655–686.
- Wang, X., T. M. Hamill, J. S. Whitaker, and C. H. Bishop, 2007: A comparison of hybrid ensemble transform kalman filter optimum interpolation and ensemble square root filter analysis schemes. *Monthly Weather Review*, **135** (3), 1055–1076, doi:10.1175/MWR3307.1.
- Whitaker, J. S., and T. M. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Monthly Weather Review*, **130** (7), 1913–1924, doi:10.1175/1520-0493(2002)130<1913:EDAWPO>2.0.CO;2.
- Würsch, M., and G. C. Craig, 2014: A simple dynamical model of cumulus convection for data assimilation research. *Meteorologische Zeitschrift*, **23** (5), 483–490, doi:10.1127/0941-2948/2014/0492.

LIST OF TABLES

Table 1. Mean, median, 10% and 90% quantile of RMSE averaged over 50 trials. The number in the parentheses is the summary statistics' corresponding standard deviation. 18

TABLE 1: Mean, median, 10% and 90% quantile of RMSE averaged over 50 trials. The number in the parentheses is the summary statistics' corresponding standard deviation.

$n = 400$	10%	50 %	Mean	90%
TAPER-EnKF	0.580 (.01)	0.815 (.01)	0.878 (.02)	1.240 (.03)
PEnKF	0.538 (.02)	0.757 (.03)	0.827 (.03)	1.180 (.05)
$n = 100$	10%	50 %	Mean	90%
TAPER-EnKF	0.582 (.01)	0.839 (.02)	0.937 (.03)	1.390 (.06)
PEnKF	0.717 (.04)	0.988 (.04)	1.067 (.04)	1.508 (.05)
$n = 25$	10%	50 %	Mean	90%
TAPER-EnKF	0.769 (.04)	1.668 (.13)	1.882 (.09)	3.315 (.11)
PEnKF	0.971 (.03)	1.361 (.03)	1.442 (.03)	2.026 (.04)
$n = 10$	10%	50 %	Mean	90%
TAPER-EnKF	2.659 (.07)	3.909 (.06)	3.961 (.05)	5.312 (.06)
PEnKF	1.147 (.02)	1.656 (.02)	1.735 (.02)	2.437 (.04)

LIST OF FIGURES

- Fig. 1.** Each line represents the normalized values of entries of row i of the inverse covariance matrix, ordered from $i - 20$ to $i + 20$, where the values are averaged over 50 trials. The PEnKF algorithm is successful at identifying that the state variables far away from variable i have no effect on it, even though there are fewer ensemble members than state variables. . . . 20
- Fig. 2.** The RMSE of the TAPER-EnKF and PEnKF over 50 trials. The darker lines of each linetype are the mean and the colored areas are the 95% confidence intervals. There is clear separation between the RMSE of the two filters with the PEnKF's error as significantly smaller. 21
- Fig. 3.** Fluid height, rain, and wind at 300 different locations at an instance of time. The blue dots are observations; rain is always observed, wind is only observed when the rain is non-zero, fluid height is never observed. The dashed lines in fluid height are the cloud and rainwater thresholds. . . . 22
- Fig. 4.** The RMSE of the TAPER-EnKF, NAIVE-LEnKPF, and PEnKF over 50 trials. The darker lines of each linetype are the mean and the colored areas are the 95% confidence intervals. All three filters are pretty indistinguishable except for the PEnKF's rain error, which is statistically smaller than the others. . . . 23




FIG. 1: Each line represents the normalized values of entries of row i of the inverse covariance matrix, ordered from $i - 20$ to $i + 20$, where the values are averaged over 50 trials. The PEnKF algorithm is successful at identifying that the state variables far away from variable i have no effect on it, even though there are fewer ensemble members than state variables.




FIG. 2: The RMSE of the TAPER-EnKF and PEnKF over 50 trials. The darker lines of each linetype are the mean and the colored areas are the 95% confidence intervals. There is clear separation between the RMSE of the two filters with the PEnKF's error as significantly smaller.

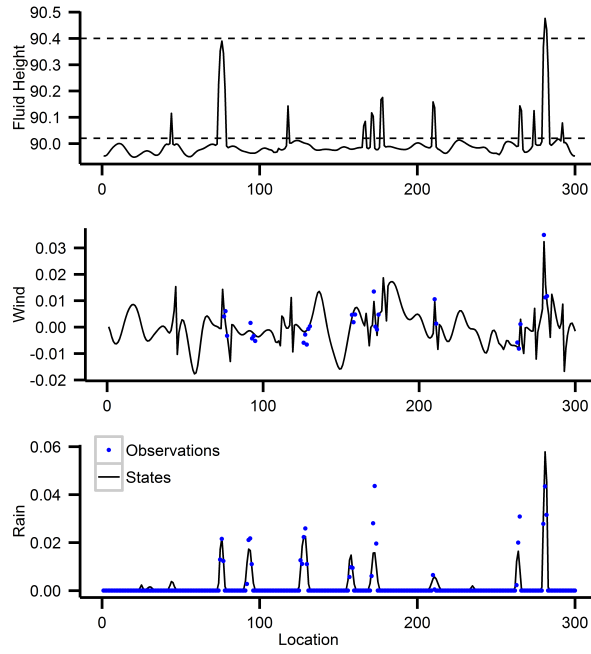


FIG. 3: Fluid height, rain, and wind at 300 different locations at an instance of time. The blue dots are observations; rain is always observed, wind is only observed when the rain is non-zero, fluid height is never observed. The dashed lines in fluid height are the cloud and rainwater thresholds.




FIG. 4: The RMSE of the TAPER-EnKF, NAIVE-LEnKPF, and PEnKF over 50 trials. The darker lines of each linetype are the mean and the colored areas are the 95% confidence intervals. All three filters are pretty indistinguishable except for the PEnKF's rain error, which is statistically smaller than the others.